

PLUS Search Results for S/N 10797650, Searched July 21, 2005

The Patent Linguistics Utility System (PLUS) is a USPTO automated search system for U.S. Patents from 1971 to the present. PLUS is a query-by-example search system which produces a list of patents that are most closely related linguistically to the application searched. This search was prepared by the staff of the Scientific and Technical Information Center, SIRA.

5838895	5619681
4616337	5623623
5398325	5663887
6219583	5826099
6219583	5826093
4534013	5825711
5581727	5828825
5666514	5929928
5778433	5946260
5831861	5953284
5893155	6026050
6115791	6061749
4298928	6061749
RE31790	6064600
4510582	6105154
4531199	6118724
4593375	6131131
4602331	6148381
4811284	6148348
5317501	6172626
5377338	6175905
5467358	6178488
5488709	6202119
5613152	6216217
5668967	6216217
5732241	3692989
5748551	3665416
5813029	4327408
5829036	4361877
5903509	4423441
5907862	4455621
6119197	4480279
6128700	4488290
6219755	4488289
6219755	4545068
4328544	4549263
4333143	4593350
5014194	4616319
5021873	4761735
5195176	4783739
5235429	4811267
5276845	4812973
5293483	4833649
5365346	4855578
5394336	4866597
5406472	4866600
5499384	4866609
5517613	4896256
5530835	4926315
5619409	4941107



Terms used

mirroring or backup or logical data recovery method

Found 80,274 of 158,639

Sort results by

☒

[Save results to a Binder](#)

[Try an Advanced Search](#)

Display results

☒

[Search Tips](#)

[Try this search in The ACM Guide](#)

☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Management of a remote backup copy for disaster recovery](#)

Richard P. King, Nagui Halim, Hector Garcia-Molina, Christos A. Polyzois

May 1991 **ACM Transactions on Database Systems (TODS)**, Volume 16 Issue 2

Full text available: pdf(2.48 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

A remote backup database system tracks the state of a primary system, taking over transaction processing when disaster hits the primary site. The primary and backup sites are physically isolated so that failures at one site are unlikely to propagate to the other. For correctness, the execution schedule at the backup must be equivalent to that at the primary. When the primary and backup sites contain a single processor, it is easy to achieve this property. However, this is harder to do when ...

Keywords: database initialization, hot spare, hot standby, remote backup

2 [A comparison of high-availability media recovery techniques](#)

George Copeland, Tom Keller

June 1989 **ACM SIGMOD Record , Proceedings of the 1989 ACM SIGMOD international conference on Management of data**, Volume 18 Issue 2

Full text available: pdf(1.32 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We compare two high-availability techniques for recovery from media failures in database systems. Both techniques achieve high availability by having two copies of all data and indexes, so that recovery is immediate. "Mirrored declustering" spreads two copies of each relation across two identical sets of disks. "Interleaved declustering" spreads two copies of each relation across one set of disks while keeping both copies of each tuple on separate disks. Both ...

3 [Session 3: Synapse approach to database recovery](#)

Kee S. Ong

April 1984 **Proceedings of the 3rd ACM SIGACT-SIGMOD symposium on Principles of database systems**

Full text available: pdf(492.53 KB)

Additional Information: [full citation](#), [references](#), [citations](#)

4 Highly available systems for database applications

Won Kim

March 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 1

Full text available:  pdf(2.43 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

As users entrust more and more of their applications to computer systems, the need for systems that are continuously operational (24 hours per day) has become even greater. This paper presents a survey and analysis of representative architectures and techniques that have been developed for constructing highly available systems for database applications. It then proposes a design of a distributed software subsystem that can serve as a unified framework for constructing database applica ...

5 Comparison of access methods for time-evolving data

Betty Salzberg, Vassilis J. Tsotras

June 1999 **ACM Computing Surveys (CSUR)**, Volume 31 Issue 2

Full text available:  pdf(529.53 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper compares different indexing techniques proposed for supporting efficient access to temporal data. The comparison is based on a collection of important performance criteria, including the space consumed, update processing, and query time for representative queries. The comparison is based on worst-case analysis, hence no assumptions on data distribution or query frequencies are made. When a number of methods have the same asymptotic worst-case behavior, features in the methods tha ...

Keywords: I/O performance, access methods, structures, temporal databases

6 Fast cluster failover using virtual memory-mapped communication

Yuan Yuan Zhou, Peter M. Chen, Kai Li

May 1999 **Proceedings of the 13th international conference on Supercomputing**

Full text available:  pdf(1.45 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

7 The HP AutoRAID hierarchical storage system

John Wilkes, Richard Golding, Carl Staelin, Tim Sullivan

February 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 1

Full text available:  pdf(1.82 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Configuring redundant disk arrays is a black art. To configure an array properly, a system administrator must understand the details of both the array and the workload it will support. Incorrect understanding of either, or changes in the workload over time, can lead to poor performance. We present a solution to this problem: a two-level storage hierarchy implemented inside a single disk-array controller. In the upper level of this hierarchy, two copies of active data are stored to provide f ...


Keywords: RAID, disk array, storage hierarchy

8 Evaluation of remote backup algorithms for transaction-processing systems

Christos A. Polyzois, Héctor García-Molina

September 1994 **ACM Transactions on Database Systems (TODS)**, Volume 19 Issue 3

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index](#)

Full text available:  pdf(1.87 MB)

[terms](#), [review](#)

A remote backup is a copy of a primary database maintained at a geographically separate location and is used to increase data availability. Remote backup systems are typically log-based and can be classified into 2-safe and 1-safe, depending on whether transactions commit at both sites simultaneously or first commit at the primary and are later propagated to the backup. We have built an experimental database system on which we evaluated the performance of the epoch and the dependency recons ...

Keywords: disaster recovery, hot spare, hot standby, remote backup

9 [Recovery in the Calypso file system](#)

Murthy Devarakonda, Bill Kish, Ajay Mohindra

August 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 3

Full text available:  pdf(318.88 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)


This article presents the design and implementation of the recovery scheme in Calypso. Calypso is a cluster-optimized, distributed file system for UNIX clusters. As in Sprite and AFS, Calypso servers are stateful and scale well to a large number of clients. The recovery scheme in Calypso is nondisruptive, meaning that open files remain open, client modified data are saved, and in-flight operations are properly handled across server recover. The scheme uses distributed state amount the client ...

Keywords: Calypso, cluster systems, distributed state, state reconstruction

10 [Distributed file systems: concepts and examples](#)

Eliezer Levy, Abraham Silberschatz

December 1990 **ACM Computing Surveys (CSUR)**, Volume 22 Issue 4

Full text available:  pdf(5.33 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The purpose of a distributed file system (DFS) is to allow users of physically distributed computers to share data and storage resources by using a common file system. A typical configuration for a DFS is a collection of workstations and mainframes connected by a local area network (LAN). A DFS is implemented as part of the operating system of each of the connected computers. This paper establishes a viewpoint that emphasizes the dispersed structure and decentralization of both data and con ...

11 [The HP AutoRAID hierarchical storage system](#)

J. Wilkes, R. Golding, C. Staelin, T. Sullivan

December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles**, Volume 29 Issue 5

Full text available:  pdf(1.60 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

12 [Optimistic recovery in distributed systems](#)

Rob Strom, Shaula Yemini

August 1985 **ACM Transactions on Computer Systems (TOCS)**, Volume 3 Issue 3

Full text available:  pdf(1.75 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Optimistic Recovery is a new technique supporting application-independent transparent recovery from processor failures in distributed systems. In optimistic recovery

communication, computation and checkpointing proceed asynchronously. Synchronization is replaced by *causal dependency tracking*, which enables a posteriori reconstruction of a consistent distributed system state following a failure using *process rollback* and *message replay*.

Because there is no s ...

13 Performance analysis of recovery techniques

Andreas Reuter

December 1984 **ACM Transactions on Database Systems (TODS)**, Volume 9 Issue 4

Full text available:  pdf(2.47 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Various logging and recovery techniques for centralized transaction-oriented database systems under performance aspects are described and discussed. The classification of functional principles that has been developed in a companion paper is used as a terminological basis. In the main sections, a set of analytic models is introduced and evaluated in order to compare the performance characteristics of nine different recovery techniques with respect to four key parameters and a set of other pa ...

14 Understanding fault-tolerant distributed systems

Flavin Cristian

February 1991 **Communications of the ACM**, Volume 34 Issue 2


Full text available:  pdf(6.17 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

15 Fault tolerance under UNIX

Anita Borg, Wolfgang Blau, Wolfgang Graetsch, Ferdinand Herrmann, Wolfgang Oberle

January 1989 **ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 1

Full text available:  pdf(1.97 MB)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The initial design for a distributed, fault-tolerant version of UNIX based on three-way atomic message transmission was presented in an earlier paper [3]. The implementation effort then moved from Auragen Systems1 to Nixdorf Computer where it was completed. This paper describes the working system, now known as the TARGON/32. The original design left open questions in at least two areas: fault tolerance for server processes and recovery after a crash were brie ...

16 Mixed integer programming methods for computing nonmonotonic deductive databases

Colin Bell, Anil Nerode, Raymond T. Ng, V. S. Subrahmanian

November 1994 **Journal of the ACM (JACM)**, Volume 41 Issue 6

Full text available:  pdf(2.54 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Though the declarative semantics of both explicit and nonmonotonic negation in logic programs has been studied extensively, relatively little work has been done on computation and implementation of these semantics. In this paper, we study three different approaches to computing stable models of logic programs based on mixed integer linear programming methods for automated deduction introduced by R. Jeroslow. We subsequently discuss the relative efficiency of these algorithms. The results of ...

Keywords: deductive databases, logic programming, nonmonotonic reasoning, operations research

17 Special issue: AI in engineering

D. Sriram, R. Joobbani

January 1985 **ACM SIGART Bulletin**, Issue 91


Full text available:  pdf(8.79 MB) Additional Information: [full citation](#), [abstract](#)

The papers in this special issue were compiled from responses to the announcement in the July 1984 issue of the SIGART newsletter and notices posted over the ARPAnet. The interest being shown in this area is reflected in the sixty papers received from over six countries. About half the papers were received over the computer network.

18 The impact of object technology on commercial transaction processing

Edward E. Cobb

August 1997 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 6 Issue 3

Full text available:  pdf(649.17 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Businesses today are searching for information solutions that enable them to compete in the global marketplace. To minimize risk, these solutions must build on existing investments, permit the best technology to be applied to the problem, and be manageable. Object technology, with its promise of improved productivity and quality in application development, delivers these characteristics but, to date, its deployment in commercial business applications has been limited. One possible reason is the ...

Keywords: Objects, Workflow, transaction processing

19 Illustrative risks to the public in the use of computer systems and related technology

Peter G. Neumann

January 1996 **ACM SIGSOFT Software Engineering Notes**, Volume 21 Issue 1

Full text available:  pdf(2.54 MB) Additional Information: [full citation](#)

20 Ext3cow: a time-shifting file system for regulatory compliance

Zachary Peterson, Randal Burns

May 2005 **ACM Transactions on Storage (TOS)**, Volume 1 Issue 2

Full text available:  pdf(443.01 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The ext3cow file system, built on the popular ext3 file system, provides an open-source file versioning and snapshot platform for compliance with the versioning and auditability requirements of recent electronic record retention legislation. Ext3cow provides a *time-shifting* interface that permits a real-time and continuous view of data in the past. Time-shifting does not pollute the file system namespace nor require snapshots to be mounted as a separate file system. Further, ext3cow is i ...

Keywords: Versioning file systems, copy-on-write

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)